



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

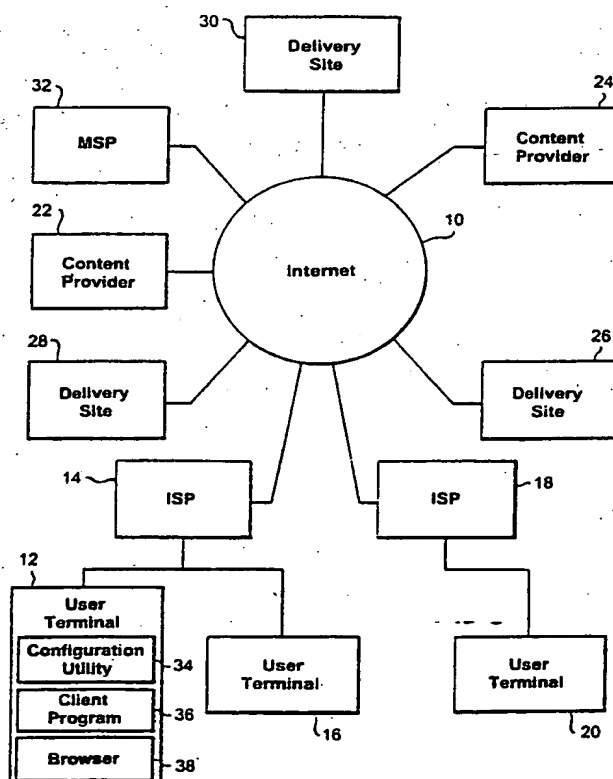
(51) International Patent Classification 6 : G06F 9/46, 17/30		A1	(11) International Publication Number: WO 98/58315
			(43) International Publication Date: 23 December 1998 (23.12.98)
(21) International Application Number: PCT/US98/12784		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 16 June 1998 (16.06.98)			
(30) Priority Data: 08/878,385 18 June 1997 (18.06.97) US			
(63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Application US 08/878,385 (CON) Filed on 18 June 1997 (18.06.97)			
(71) Applicant (for all designated States except US): INTERVU, INC. [US/US]; 6815 Flanders Drive, San Diego, CA 92121 (US).			
(72) Inventors; and (75) Inventors/Applicants (for US only): KENNER, Brian [US/US]; 1403 Walnut Creek Drive, Encinitas, CA 92024 (US). COLBY, Kenneth, W. [US/US]; 12707 Gibraltar Drive, San Diego, CA 92128 (US). MUDRY, Robert, N. [US/US]; 4350 Point Reyes Court, Carlsbad, CA 92008 (US).			
(74) Agents: SCHAFFER, Robert et al.; Darby & Darby P.C., 707 Wilshire Boulevard, Los Angeles, CA 90017-3514 (US).			

Published
With international search report.

(54) Title: SYSTEM AND METHOD FOR SERVER-SIDE OPTIMIZATION OF DATA DELIVERY ON A DISTRIBUTED COMPUTER NETWORK

(57) Abstract

A system and method for the optimized storage and retrieval of video data at distributed sites calls for the deployment of "Smart Mirror" sites throughout a network, each of which maintains a copy of certain data managed by the system. User addresses are assigned to specific delivery sites based on an analysis of network performance with respect to each of the available delivery sites. Generalized network performance data is collected and stored to facilitate the selection of additional delivery sites and to ensure the preservation of improved performance in comparison to traditional networks.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav	TM	Turkmenistan
BF	Burkina Faso	GR	Greece		Republic of Macedonia	TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

5

10

**SYSTEM AND METHOD FOR SERVER-SIDE
OPTIMIZATION OF DATA DELIVERY ON A
DISTRIBUTED COMPUTER NETWORK**

15

The invention relates to a system and method for distributed data storage and retrieval, and more particularly, to a system and method whereby a user can acquire network performance information for a dynamic and distributed multipurpose network, and this information will then be used to identify and
20 select optimum delivery sites or servers from which to receive computer data, specifically multimedia content. Such delivery sites and servers are selected so as to increase network capacity, distribute server load, and reduce transmission delays between the server and the user.

25

BACKGROUND OF THE INVENTION

The Internet is a loose network of connected computers spread throughout the world. A message can be sent from any computer on the Internet to any other by specifying a destination address and passing the message from computer to computer via a series of "hops." Each computer,
30 router, or "node" on the Internet has a unique Internet address. When an intermediate computer or router receives a message in transit, the computer checks the intended destination of the message and passes it along accordingly.

The Internet is growing, in terms of both size and sophistication, at a rapid rate. In the past, most users of the Internet were academic, research, or
35 institutional users; the Internet was primarily used at that time to transmit and

receive electronic mail and network news and to allow transfer of computer files. However, since the introduction of the World Wide Web (also known as the "Web" or the "WWW") several years ago, the Internet has begun to host increasing amounts of other types of data of general interest, namely
5 representations of images, articles, etc.

The Web protocol and language establish a graphical means to navigate the expanses of the Internet. "Web pages," often consisting primarily of text and graphical material, are stored on numerous computers, known as "Web servers," throughout the Internet. A software program known as a "browser"
10 can be used to access and view Web pages across the Internet by specifying the location (i.e. Internet address) of the desired Web page. When a Web page is accessed, its information is transmitted from the remote computer (server or delivery site), wherever in the world it may be located, across the Internet, to the user.

15 In recent times, the Web has begun to host highly sophisticated types of multimedia content, such as audio and video data, and computer software. Compared to first generation Web content, namely text and still images, audio clips, video clips, and software programs have extremely high storage and bandwidth requirements.

20 At present, it is difficult, if not impossible, to provide sustained high-speed transmission of large audio/video files over a multi-node link on the Internet. Because the data is often transferred from afar, many factors can cause the delay or even loss of parts or all of a transmission. It is generally not critical if a user experiences minor delays in receiving small graphic or text files.
25 However, it is recognized that real-time data such as video has very specific and stringent timing requirements for data transfer and display.

Unfortunately, the present design of traditional Internet-like data networks is based on the principle that delays and significant data transmission rate variations are acceptable for ordinary data (e.g. text and still images).
30 Consequently, because of the high value of permitting access to text and graphical information from locations around the world, such transmission defects are considered acceptable, and the base capacity of the Internet is somewhat

"oversubscribed" to reduce data transmission costs. In other words, the timeliness of network data transmission has been significantly compromised in order to render relatively insignificant the aggregate cost of long distance communication connections.

5 In order to successfully transfer audio-video data across a message-oriented network such as the Internet, for any more than a few users, network resources should be committed in a manner facilitating timeliness of transmittal. A system using committed network resources generally cannot take advantage of the existing pricing scheme of shared networks like the Internet, since it cannot participate in the sharing of network resources on a data packet by data
10 packet basis. Video data must be transmitted to the exclusion of lower-priority data. Transmission costs thus become significant, especially when the connection is "long distance" or when the connection is continued over an extended period of time.

15 Another consequence of the timeliness vs. cost compromise discussed above has been the seemingly indiscriminate topographical design of the network. Since delays and throughput variations have traditionally been excused in favor of low cost, the configuration of the Internet infrastructure has also been driven by cost considerations. Accordingly, the interconnection efficiency
20 of the network has rarely been considered. The rapid growth of real time data is changing this requirement.

It is recognized that inadequate data transfer performance of time-sensitive data on the Internet is typically caused by four factors: packet loss, excessive server utilization, the relatively low capacity of the network
25 infrastructure, and inherent delays in the network hardware. Packet loss, in particular, is caused by inadequate infrastructure and lack of robustness in routing. The inherent delays are believed to be caused by, among other things, the lack of flow control between adjacent nodes in a multiple-node path on the Internet.

30 Unlike smaller text and graphic files, relatively large video files can take several minutes (or more) of "streaming," or constant data flow. Consequently, the usual network performance problems are exacerbated. Network bandwidth,

or the data-carrying capacity of a particular network, is limited. Thus, packet loss and delays increase. Long delivery times consume a large amount of server capacity for a long time, decreasing the resources available to other users. Accordingly, because the network infrastructure becomes increasingly congested, packet loss and delays continue to increase, transmission times rise, and server load increases further.

This pattern exemplifies a "downward spiral" of network performance, which can be driven by the attempted transmission of large data files such as video clips. As long as network traffic remains within the limits imposed by network bandwidth, network performance will remain acceptable. However, whenever peak network loads exceed capacity, the downward spiral described above will begin, causing increasing periods of poor network performance.

As discussed above, a browser program can be used to access and view Web pages across the Internet by specifying the location (i.e. Internet address) of the desired Web page, or more commonly, by "hotlinking" to Web pages. Common browsers are Lynx, NCSA Mosaic, Netscape Navigator, and Microsoft Internet Explorer. The desired Web page is specified by a uniform resource locator ("URL"), indicating the precise location of the file using the syntax "http://internet.address/directory/filename.html".

Web pages are generally described, in terms of layout and content, by way of a language known as "HTML" (HyperText Markup Language). Any particular computer linked to the Internet can store one or more Web pages, i.e. computer files in HTML format, for access by users.

Hotlinking from one HTML Web page to another is accomplished as follows. The user first accesses a Web page having a known address, often on the computer located at the user's ISP (Internet Service Provider). The ISP is the organization providing Internet connectivity to the user. That Web page can contain, in addition to textual and visual data specified in HTML format, "links," or embedded information (in the form of URLs) pointing to the Internet addresses of other Web pages, often on other computers throughout the Internet. The user, by selecting a link (often by pointing and clicking with a mouse), can then access other Web pages, which can in turn contain further data and/or additional

links.

Various extensions to HTML, such as Netscape's EMBED tag, allow references to other data to be embedded into Web pages. Some browsers are not capable of handling data other than text and images. Other browsers can
5 handle the data in various ways. NCSA Mosaic, for example, handles references to unknown types of data by allowing the data to be downloaded to the user's computer, and then optionally invoking an external program to view or manipulate the data. Recent releases of Netscape Navigator and Microsoft Internet Explorer take the concept one step further: a browser extension, or
10 "plug-in," can be automatically invoked to handle the data as it is received from the remote Web page. Other means, such as network program "applets" written in the Java language (or a similar language), can be used to extend the functionality of the browser environment or network.

Digital multimedia data can have extremely high storage and bandwidth
15 requirements. In particular, video files can be very large, from approximately 10 megabytes to 10 gigabytes. In order to play video files at speeds approaching their recorded rate at a user's terminal, the files have to be delivered at a fast, constant speed. Too slow, and the image plays back slower than originally recorded. If the speed is uneven, then the video appears jerky, like an old-time
20 movie.

The network design compromises discussed above generally adversely impact the transmission of audio and video data across the Internet. While a user using a browser to "surf" the Web might not notice minor delays and transmission rate variations while retrieving text and still images, such defects
25 become apparent and significant when real-time audio and video information is accessed.

In an attempt to solve these problems, Internet content providers sometimes spread popular content around the Internet on various servers or delivery sites known as "mirror sites." Each mirror site contains information that
30 is essentially identical to that of the original site. For example, if a popular Web site is located in New York, mirror sites might be located in Los Angeles, London, and Tokyo. Accordingly, if a European user is having difficulty

accessing the original New York site, he can hotlink to the mirror site that is geographically closest, i.e. London.

However, mirror sites have several disadvantages. For example, mirror sites may be widely distributed geographically, but may not be efficiently distributed on the network in terms of actual usage, network traffic, etc. Thus, New York and Los Angeles mirror sites might both be connected to the same national Internet service provider's network, meaning that difficulty in accessing one of the sites might also affect the other.

Furthermore, mirror sites might not be optimally placed to reduce load on each server. Although an "educated guess" might be made as to where a mirror site should be located, actual usage patterns might differ. Furthermore, there is no guarantee of enhanced performance. The bandwidth of the mirror site might be lower than that of the original site, or it might be overloaded for other reasons.

Moreover, mirror sites are often hosted on a voluntary basis. If a Web site is extremely popular, and a service provider determines that the subject matter might be of interest to its subscribers, that service provider might agree to host a mirror site of the original Web site. Such an arrangement would be attractive to host of the mirror site because people would be drawn to the mirror site, and might hotlink to other content hosted there. On the other hand, such voluntary alliances typically are not reliable and might be severed at any time.

In essence, a mirror site offers a secondary source for data, which may or may not be available, and which may improve user convenience, but which does not address network bandwidth or efficiency. A mirror site does not account for performance characteristics of the network, nor identify available bandwidth which could be used to efficiently transmit video data while still taking advantage of the existing low-cost pricing schemes such as those on the Internet.

Currently, there is no guidance in selecting optimal locations for delivery sites, nor is there a known method permitting a user to determine which mirror site to connect to that will ensure optimum performance. In fact, the use of a traditional mirror site is voluntary. Typically, a user will try to access the original

site (or a known mirror site), and will switch to another mirror site only if performance is found to be insufficient after one or more attempts. This approach is an inefficient utilization of network resources. Clearly, mirror sites are not an optimum solution to the problem of overloaded Web sites. A principal
5 reason for this, among others, is the failure to consider network performance.

Network analysis, particularly the performance of specific paths and links over the Internet, is well known and developed. For example, the "ping" program allows a computer connected to the Internet to determine whether a remote host is accessible. However, the ping program uses a low-priority
10 network protocol known as the ICMP protocol, and accordingly does not provide meaningful performance analysis information. The "traceroute" program follows the transmission of a message from a computer to a remote host, tracking delays along each link, and determining the path taken by the message. The traceroute application can be used to map the flow of data. However, it lacks
15 the ability to provide meaningful performance analysis information. Traceroute only provides route information for a message propagating in one direction, and only for one instant in time.

Moreover, only the connectivity characteristics of paths leading to and from the single computer running the tests are typically determined; expanding
20 the scope of testing is possible but logistically impracticable, since the Internet is so large.

Traditional network analysis techniques such as the "ping" and "traceroute" programs offer a view of network connectivity but provide little understanding of what performance can be expected from providers and mirror
25 sites across the Internet. Therefore, only "guesses" can be made as to where delivery or mirror sites should be located or which mirror sites should be used to optimize performance.

Accordingly, a need exists for a method of determining overall network performance. A further need exists for a system applying that method to enable
30 content providers to dynamically locate data delivery or mirror sites at optimum network locations, and to allow users to select optimum mirror sites from which to receive data.

SUMMARY OF THE INVENTION

The invention is directed to a system and method for the optimized distribution of Web content to sites located around the Internet. An intelligent mirroring scheme, called here "Smart Mirroring," is used to determine the need
5 for and distribution of mirror sites and to direct user requests for certain Web content to an optimum mirror site.

A number of "smart" delivery or mirror sites are used to distribute popular Web content to various parts of the Internet. A comprehensive scheme of network analysis, based on tests performed by a large number of users, is used
10 to interactively determine the preferred locations for the sites, and to determine the optimum sites to be used by each individual user.

Accordingly, because each individual user is routed to a Smart Mirror or delivery site that provides improved performance, overall network congestion is reduced. In most cases, the improved server is located electronically close to
15 a user in order to decrease the number of network connections over which data must travel, thereby reducing packet loss and delay.

Furthermore, network analysis results allow message traffic to be routed away from those delivery sites and network regions that are already overloaded, and toward underutilized servers and networks. This results in an improvement
20 in throughput as seen by each user, and will thereby increase the appeal of the content offered by content providers using the system. Content providers are able to reach a larger number of users across the Internet without suffering significant decreases in performance.

A system according to the invention begins with an original Web site and
25 at least one additional delivery (or mirror) site. Each user desiring to use the system will be provided, in a preferred embodiment, with software which includes a configuration utility and a client program. The configuration utility is used first to determine which delivery sites provide improved performance for that particular user.

30 In one embodiment of the invention, the configuration utility first downloads a "delivery site file" from a service provider. This delivery site file contains a list of available delivery sites and a list of network tests to be run.

The types of tests and frequency of testing to be performed may be specified in the delivery site file, as dependent on the number of users testing the network and the estimated drain on network or delivery system capacity.

The configuration utility will run a subset of the tests specified in the delivery site file. The test results show which delivery sites yield improved performance for the user, and also contain information on various generalized network capabilities from the standpoint of the user running the tests. The network test results and the identity of the chosen delivery site will be sent (via e-mail in one possible configuration) back to the delivery service provider for incorporation into the service provider's database.

The delivery site chosen by the configuration utility is then used by that user for the retrieval of all content managed by the delivery system service provider. Consequently, when the user is browsing Web content, and finds a particular item, e.g. a video clip, that is managed by the service provider's delivery system, the client software will automatically retrieve it from the specified "Smart Mirror" delivery site. Site preferences and default sites can be updated periodically on request, at specified times, or in response to changes in network load and traffic.

Moreover, because the configuration utility of the invention is performing various network tests and providing the test results to the service provider, valuable data on system and network performance is available. Such data provides information on which "Smart Mirror" delivery sites are performing effectively and which are not, which Smart Mirror delivery sites are overloaded, and what portions of the Internet might benefit from the addition of more delivery sites or capacity. Such data also makes it possible to perform such sophisticated network analysis as end-to-end performance measurements, workload characterization, route stability, and outage metrics.

In an embodiment of the invention, the mirror service provider uses the network performance data provided by the end users to derive a look-up table which correlates Internet IP addresses with "electronically close" delivery sites. When a user is browsing web pages and requests a file, e.g. an advertising banner or video clip that is managed by the service provider's delivery system,

the service provider can map the user's IP address to the look-up table and determine which delivery sites are "electronically close" to the user. The service provider can then provide the user's configuration utility or client program with a single delivery site address or a list of delivery site addresses for these servers.

- 5 In the latter case, the user terminal acts as a router making the final delivery site selection.

In general, an improved delivery site for a particular user can be predicted in advance by analyzing aggregate network performance data collected from network tests previously performed by a group of users. Thus, delivery site
10 selection can occur on-the-fly each time the user requests a file managed by the mirror service provider's delivery system. From the perspective of the user, the selection of the delivery site happens automatically and transparently such that there appears to be no delay between selecting a file from a web page and having the file delivered to the user's terminal. The look-up list maintained by
15 the service provider is constantly updated to reflect changes in network performance, making it possible for the service provider to effectively load-balance network traffic.

Thus, from an engineering standpoint, the mirror service provider can continue to ensure that improved performance is being provided. From a
20 marketing perspective, content providers can be told where to locate Smart Mirror or delivery sites for improved performance, and what ISP provides improved delivery.

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 is a block diagram of an illustrative network topology of a system according to the invention, including multiple users and multiple content providers;

5 FIGURE 2 is a flowchart describing the operation of the configuration utility used in a system according to the invention;

FIGURE 3 is a flowchart describing the operation of a client program used in a system according to the invention; and

10 FIGURE 4 is a flowchart illustrating how site selection is performed in an embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

The invention is described below, with reference to detailed illustrative embodiments. It will be apparent that the invention can be embodied in a wide
15 variety of forms, some of which may be quite different from those of the disclosed embodiments. Consequently, the specific structural and functional details disclosed herein are merely representative and do not limit the scope of the invention.

Referring initially to FIG. 1, the Internet 10, which is intended to be
20 representative of wide-area communications networks in general, is depicted as a "cloud." The Internet is known to be an interconnected network of a large number of computers. Although Internet-connected computers that are "geographically" near each other can be "electronically" near each other on the Internet, such is not usually the case. However, one computer connected to the
25 Internet can communicate with any other computer connected to the Internet; the message will most likely travel over a path comprising a sequence of links, or "hops," between computers that are directly connected to each other.

A first user terminal 12 is also depicted in FIG. 1. The first user terminal 12 is connected to an Internet service provider (ISP) 14, which is typically just
30 a computer, router, or terminal server connected to the Internet 10. An ISP 14 can host additional user terminals, such as a second user terminal 16. Other ISPs, such as a second ISP 18, are also connected to the Internet 10. A third

user terminal 20 is shown connected to the second ISP 18. Only three user terminals are shown; however, it should be recognized that the number of concurrent users of the invention is unlimited, subject to the operational details set forth below.

5 As is known in the art, content providers are also connected to the Internet 10. A first content provider 22 might provide a certain kind of content, for example sports scores and highlights. A second content provider 24 might provide a different kind of content, for example business news.

10 Traditionally, if a user (such as the one using the first user terminal 12) wished to access the content provided by the first content provider 22, the terminal 12 would query the first content provider 22 directly. A request message would propagate from the terminal 12, across the Internet 10, to the content provider 22. The content provider 22 would send the desired data across the Internet 10 back to the terminal 12.

15 Several delivery, or "mirror" sites are shown connected to the Internet 10 in FIG. 1. A first delivery site 26 might be located a small number of "hops" from the first user terminal 12. A second delivery site 28 might be located further away from the first user terminal 12, but close to the third user terminal 20. A third delivery site 30 might be as close to the third user terminal 20 as the second delivery site 28 is. As previously noted, a user and a provider or delivery site that are "geographically" near each other might not be "electronically" near each other on the Internet. By decreasing the "electronic" distance between the user and the provider or delivery site, the number of network connections and routers over which data must travel can be decreased.

25 As discussed above, the Smart Mirroring system acts to improve network performance by decreasing the incidence of the foregoing network problems. Packet loss and delay problems are generally decreased by reducing the number of network connections over which data must travel, although in some cases, the network testing procedure of the invention shows that some longer paths provide better throughput than some shorter paths. Very little packet loss, and essentially no delay, occurs in network cable; it typically is caused by overloaded network storage and routing devices. Because the Smart Mirror sites of the

30

invention are located electronically near each user, packet losses and delays are reduced. The problem of excessive server utilization is reduced because multiple delivery sites share the load that typically would have been handled by a single server. The relatively low capacity of the network infrastructure becomes less of a problem, because data retrieved from parallel delivery sites in different locations generally need not travel over the same network links.

For the purposes of describing this invention, a delivery site is a "node" on the network which may store data or other files, such as software code, for delivery. The term can also include a site which is responsible for data delivery, including mirror sites, content providers, and servers for broadcast video streams or Web sites.

In the system, a mirror service provider (MSP) 32 is connected to the Internet 10. The MSP 32, which exercises a management function over the distribution of delivery sites 26, 28, and 30, and over the allocation of requests to the original and delivery sites from user terminals 12, 16, and 20, includes a database capable of transmitting and receiving data over the Internet 10.

This management function is facilitated by the use of a configuration utility 34 and a client program 36 run within a storage medium (i.e. random access memory) on the user terminal 12. Although the configuration utility 34 and the client program 36 are shown in FIG. 1 as a part of only the first user terminal 12, it should be recognized that any user terminal, such as terminals 16 and 20, participating in the system will use such software. A user desiring to participate in the system can obtain the software comprising the configuration utility 34 and client program 36 directly from the MSP 32, or through traditional retail or other channels (such as being part of the browser or operating system of the computer). It should be noted that the functions performed by the configuration utility 34 in the described embodiment of the invention can be integrated into general Internet application software, such as a browser or other network application; a stand-alone program is not necessary.

In a preferred embodiment, the configuration utility 34 must be run by the user, either by command or automatically, before the user terminal 12 will have access to the system. The operation of the configuration utility 34 is shown in

detail in FIGURE 2.

The configuration utility 34, when first run on the user terminal 12, retrieves a delivery site file (step 40) from the MSP 32 (FIGURE 1). If the user already has a delivery site file (e.g., it was received with the configuration utility 34), and that delivery site file is sufficiently new, the delivery site file can be retrieved from the local hard disk of the user terminal 12. This delivery site file contains a list of all available delivery sites (such as delivery sites 26, 28, and 30) and a list of network tests to be run at the user terminal 12. In the context of the invention, there can be as few as two delivery sites, or if the number of users justifies it, as many as several thousand. The number of sites in principal is unlimited, with each available delivery site represented in the delivery site file.

The delivery site file is generated by the database from within the MSP's computer system. The database application uses information about the user to dynamically determine the optimum tests to run. Consequently, the delivery site file need not contain entries for every delivery site in existence; the list can be tailored to include only those sites which appear appropriate or feasible.

Initially, the magnitude of run-time variation in test configurations for the delivery system users is low; that is, the first group of users all run essentially the same tests. As the delivery service grows, however, the intensity of each user's testing is reduced in order to compensate for the increased magnitude of testing network-wide. The scope of testing and the number of delivery sites tested both can be narrowed to further reduce the aggregate load of network testing.

In one embodiment, the delivery site file will have a format generally as follows:

1. File Revision Number and Message. The file includes this field to determine whether a new version of the configuration utility 34 is available. If the revision number in the delivery site file is higher than the version number for the configuration utility, configuration is not allowed. Instead, the user would be prompted to acquire a newer version of the configuration utility 34. File revision verification as described herein ensures that the most up-to-date delivery site selection algorithms are

applied to the test data generated by the configuration utility 34.

2. A list of available Smart Mirror delivery sites. For each available delivery site, the following information is provided:

a. Host name. In the known Internet format of "www.server.com."

b. IP Address. A numerical Internet address in the known format. The address is presently a 32-bit number of the form w.x.y.z., where w, x, y, and z are each in the range of 0 to 255.

c. Alternate Name. An informal name such as "The First Mirror Site."

d. A list of tests to be executed. For each test, the following information is provided:

i. Test ID. Each type of test has a unique identifier known to the configuration utility 34.

ii. Weighting factor. Each test will be weighted by a specified percentage.

iii. Frequency. Each test is not necessarily run every time. This field specifies a probability, determining how often a particular test will be run.

iv. Additional Information (optional). For certain tests, additional information may be needed.

e. Site Preference Level. Each site can be given a weighting, or preference level, between, for example, 1 and 100. As discussed below, aggregate data in the MSP's database is used to perform network usage analysis not possible with only the single user's instantaneous end to end testing. The weighting factor provided here is used to incorporate the test results received from the service provider's database. This weighting factor is also used to limit assignment of new users to a delivery site once a predetermined maximum usage level has been reached.

f. Test Site Flag. If this flag is enabled, the foregoing tests will be run, but the site will not be assigned as a delivery site even

if it yields the best performance.

g. Content Provider Groups. Each site can belong to one or more content provider groups, thereby mirroring only certain content. If a user is not interested in the types of data hosted by a particular delivery site, then it does not need to be tested.

The configuration utility 34 then queries the user (step 42) for various items of information needed in the configuration process, for example, the user's name, e-mail address, password, modem speed, and information related to access control (e.g. what levels of various attributes are viewable by the user). The access control mechanism will be discussed in further detail below. In one embodiment of the invention, the information received from the user is encrypted and stored in a configuration file on the user terminal 12.

The configuration utility 34 then determines whether the user terminal 12 is connected to the Internet (step 42). If not, it will initiate a connection (step 44) or prompt the user to do so.

A series of network tests is then performed (step 46). One or more tests can be performed for each available site listed in the delivery site file; not all sites in the file need to be tested.

The following test types are presently considered to provide useful data:

1. Ping. Provides information on whether a remote server is reachable, and if so, how long it takes for a low-priority message to travel round trip from the user terminal 12 to the remote server and back. Ping is a simple test useful in deciding whether a site is available for further evaluation. Excessive times returned by the ping application can be used to eliminate delivery systems which are far too "slow" for effective information delivery. This test is used by the terminal to reduce the number of delivery sites tested.

2. Traceroute. Provides information on what route is taken by a message from the user terminal 12 to a remote server, including what systems are used along the way, and how long each hop takes. Traceroute is used by the configuration program 34 to document the path

of information transmission. Several traces with differing results might indicate that the stability of the route from a particular user to a specific server is not acceptable. Previously aggregated data on particular routes, from the service provider's system database, may also influence the decision to choose a particular delivery site for a specific user. Route stability is the primary consideration.

3. Reverse Traceroute. Provides information on what route is taken by a message from a remote server to the user terminal, including what systems are used along the way, and how long each hop takes. Reverse Traceroute is used by the configuration program to document the path of information receipt. Several traces with differing results might indicate that the stability of the route from a particular server to a specific user is not acceptable. Previously aggregated data on particular routes, from the service provider's system database, may also influence the decision to choose a particular delivery site for a specific user. Again, route stability is the primary consideration.

4. Dynamic Traceroute. Similar to traceroute or reverse traceroute, but between any specified pair of computers on the Internet. Dynamic Traceroute is used by the configuration program to document a path of information transmission. Several traces with differing results might indicate that the stability of the route between two network locations is not acceptable. Previously aggregated data on particular routes, from the service provider's system database, may also influence the decision to choose a particular delivery site for a specific user. As above, route stability is the primary consideration.

5. Name Server Resolution Delay. If the numeric Internet address is unspecified, a name server lookup is performed to determine what numeric address corresponds to the desired host name. This process can take a substantial amount of time.

6. Throughput. A sample file is downloaded, or partially downloaded, from the remote server to determine the actual throughput in bytes per second.

7. Throughput variation. A sample file is downloaded, or partially downloaded, from the remote server to determine if the throughput is relatively constant or fluctuating.

8. Error rate. A sample file is downloaded, or partially downloaded, from the remote server to determine if the transmission is subject to transmission errors. This information is obtained by counting the number of error message packets returned to the user, thereby learning of the error rate from the user to the server, and by tracking the number of bytes received as a fraction of the number of bytes transferred, thereby learning of the error rate from the server to the user.

9. Packet Fragmentation. A sample file is downloaded, or partially downloaded, from the remote server to determine if the transmission is subject to fragmentation or out-of-order packet reception.

10. Capacity Query. If the remote server is so enabled, the configuration utility 34 queries the server to determine its transmission capacity and its average load. This information is collected via the Simple Network Management Protocol ("SNMP"), which is supported by nearly all Internet servers.

11. Macroscopic Network Analysis. The data accumulated by the MSP database offers a global view of network behavior. This information permits the Smart Mirror system user to have a historical view of the performance of the available delivery sites. The accumulated data is manipulated by the delivery system database to analyze network performance, in order to emphasize usage in high-capacity areas of the network, while deemphasizing usage in areas already experiencing diminished performance.

Information on how each of the foregoing tests are performed is well known in the art of network analysis. See, e.g., Bob Quinn and Dave Shute, Windows Sockets Network Programming (Addison-Wesley 1996). In one embodiment of the invention, testing is accomplished by performing a "ping" test to verify whether a server is reachable, a series of small (e.g. 20K)

downloads, a series of large (e.g. 200K) downloads, and "traceroute" and "reverse traceroute" tests to document delivery paths.

The traceroute information is used by the MSP 32 to correlate test data to information in its database; in that way, particularly bad network links and servers can be identified. Such information is provided in the delivery site file
5 discussed above; if a particular link or server is known to be unreliable, based on information obtained from other users, an individual user can be routed away from it, even if a single test gives good results.

The short downloads are used to determine server capacity. The name
10 server resolution delay can be determined by such a test, as well as the time it takes a server to begin sending data. The later result is strongly related to server load, capacity, and performance.

The long downloads allow the configuration utility 34 to determine how packet loss, network congestion, and server utilization affect file delivery. It is
15 not ideal to determine which of the foregoing factors is causing decreased performance based on test results from a single user. However, such results in the aggregate, as stored in the database maintained by the MSP 32, indicate the root causes.

It should be noted that some of the test results may be used in
20 conjunction with other test results. For example, the load on a delivery site as determined through a capacity query can be divided by the results of a throughput test to derive an average expected download time, given the load characteristics of the server.

After all specified tests are run, the results are collected and processed
25 (step 48). It is possible that certain tests were not able to be successfully performed; in such cases, the results should indicate an appropriate worst-case value (e.g. zero throughput or extremely high delay).

It is important to note that not all possible tests will be performed each time the configuration utility 34 is run. When a large number of users is using
30 the system, a substantial drain on server and network capacity would be caused by the testing procedure alone, increasing the downward spiral of network performance previously discussed.

As noted above, a test frequency number is stored in the delivery site file for the purpose of dynamically controlling the number of users performing a test. The testing performed by the configuration utility 34 is performed to achieve statistical confidence in deciding which delivery site is best suited for data
5 delivery to a particular user. Statistical confidence is obtained by either testing a small sample of users sufficiently well and using that data to influence the choice of a delivery site, or by having a large number of users each "lightly" test several available sites and using that data in the aggregate.

Accordingly, when the system is used initially, a relatively small number
10 of users are "enrolled" in the system. The delivery site file maintained by the MSP 32 reflects those conditions, and requires each user to test the network (through the configuration utility 34) relatively heavily. As the number of users increases, the delivery site file is modified to decrease the tests performed by each user. By the time a very large number of users are using the system, the
15 configuration utility 34 may predominantly test for delivery site reachability (via a "ping"-type test, as discussed above), and rely primarily on test data provided by other users and stored in the database maintained by the MSP 32. However, even when many users are testing the system, a small number of users (e.g. one in 5,000) may be selected to run a comprehensive set of tests.

20 Preferably, testing should not contribute more than approximately 5% of total server load. One way to reach this goal is to lightly test a large number of servers, yielding a group of delivery sites having adequate performance. This group of delivery sites can then be used in rotation to retrieve data. Information on multimedia clip actual download times for each of the delivery sites in the
25 group is accumulated as discussed below, and further information on delivery site performance can then be furnished to the MSP 32 transparently, without the need for further outright testing.

Accordingly, on the basis of the collected test results, and on information provided in the delivery site file by the MSP 32, the configuration utility 34
30 determines which delivery site, or group of delivery sites, is best for the user terminal 12 (step 50). This determination can be made numerically by weighting the various tests performed and comparing the score for each site.

In a presently preferred embodiment, for use in a low-traffic environment with a relatively small number of delivery sites and users, the configuration utility 34 relies primarily on ping and throughput tests for each available delivery site. An initial ping test is performed to determine if a delivery site is reachable. Short and long downloads are performed in the throughput tests to determine the maximum and minimum throughputs from the delivery site, and to determine whether throughput variation is small enough to accommodate the transmission of video data. Accordingly, those tests are all given high weights. Other tests, such as traceroute, can be performed, and the results reported to the MSP 32, without playing a role in the choice of delivery sites (such other tests can be given weights of zero, or nearly zero, for example). As the size of the system increases, and additional delivery sites and users are enrolled, the site selection formula can be altered, by changing the delivery site file contents, to reflect changing network conditions.

In one embodiment of the invention, a proprietary graphical interface is provided so that the location of the user and the locations (both geographic and electronic) of each site tested can be displayed on a monitor connected to the user terminal 12, allowing a visual indication of the relative distances between sites. In one embodiment, the display is shown in the form of a "radar screen," upon which the user terminal 12 and delivery sites are displayed as "blips" superimposed over a map of the pertinent geographical region. In order to encourage the user to use the application and to offer more network-wide data, the user interface can allow the user to enter an "ad-hoc" test site for additional performance testing. In this case, the configuration utility will test either the default Web page file (e.g., "index.html") or a specific file requested by the user. Analysis results from the user-selected site are adjusted so that reasonable comparisons can be made with results from other sites.

It should be noted that multiple sets of delivery sites can be maintained by the invention, to accommodate several groups of content providers. Each content provider might be mirrored only at certain sites. Accordingly, for each content provider having a unique set of delivery sites, a primary delivery site is selected by the configuration utility 34. To accomplish this, the foregoing tests

can be run once, and then, in one embodiment, a numerical weighting can be applied to each appropriate set of delivery sites. A plurality of Smart Mirror sites is selected, one for each content provider group. The content provider groups is specified in the delivery site file; each possible delivery site is identified as belonging to one or more content provider groups. When content provider groups are used, there can be as few as two groups; the maximum number is essentially unlimited.

It should also be noted that a prioritized ranking of delivery sites can also be generated and maintained. If this is done, failure of the primary Smart Mirror site to respond will allow the system to fall back to the next-highest ranked Smart Mirror site.

After a Smart Mirror site is selected, certain data will be sent to the MSP 32 (step 52) via e-mail or other Internet electronic protocol. The information received by querying the user, the identity of the selected Smart Mirror site, and all raw test data and results, including the time and date at which each test was run, is compiled into a text file (which is encrypted in one embodiment). Upon receipt by the MSP 32, the data is stored in a database for use in managing and analyzing the system.

Finally, the configuration utility 34 saves the identity of the selected Smart Mirror site for each set of delivery sites, or the prioritized list, to the (encrypted) configuration file (step 54). The configuration utility may also save information on relative performance for each tested delivery site. The client program 36 uses the encrypted configuration file to download data files (video clips or other content) from the appropriate Smart Mirror site.

It should be noted that in the operation of the system, the MSP 32 performs certain functions. The MSP 32 maintains the delivery site list, adding and deleting sites as necessary. The MSP 32 also maintains the database of network performance, containing information received via e-mail or other means from users running the configuration utility 34. As large amounts of data are received from numerous users, the database can provide valuable information on the performance and other characteristics of the Internet and portions thereof. Various data processing techniques are known to derive such information.

The locations of the delivery sites used with the invention are ultimately determined by a number of factors, including marketing considerations and cost/benefit analyses. However, the data stored in the MSP's database can confirm the utility of placing a delivery site at a given location on the Internet or other network. In one embodiment, servers are located on each major backbone (portion of the Internet maintained by a single corporation) and on other Internet lines serving large numbers of users, such as the major lines operated by the Regional Bell Operating Companies ("RBOCs"). In certain networks serving large numbers of users or having heavy video delivery traffic, servers can be placed at major Points of Presence ("POPs") for the network to ensure that each user has access to a fast server.

Once the configuration utility 34 has been run, the user can use the system to enable and facilitate the receipt of data files, specifically video clips, audio clips, software programs, and other content.

As time passes and the usage patterns of a user's region of the Internet change, the user might become dissatisfied with the performance of the Smart Mirror site associated with his user terminal 12. If that happens, the user is free to re-run the configuration utility 34. By that time, additional delivery sites might have been placed into service, or a different pre-existing site might perform better than the one previously assigned. Furthermore, if the player program 36 determines that the selected Smart Mirror site is not performing adequately (e.g., it has failed three times out of ten attempts), the player program 36 can prompt the user to re-run the configuration utility 34. In other embodiment of the invention the testing and mirror assignment is run automatically with each request for a file on the Smart Mirror service or at some intermittent times such as after every other, every third, every tenth or every one hundredth request.

In one embodiment of the invention, the Smart Mirror system is used to locate a delivery site from which to download a video or audio clip ("clip") referenced on a Web page. In this embodiment, the client program can be referred to or considered a "player program." The player program, in addition to carrying out the functions of the client program 36, enables the retrieval and playback of video data. Ordinarily, a browser program 38 is run on the user

terminal 12 to view Web content. Browser programs typically used include NCSA Mosaic, Netscape Navigator, and Microsoft Internet Explorer. The browser program 38 allows the user to hotlink among various Web sites on the Internet.

5 The EMBED tag is used within HTML documents to indicate which Web pages include content managed by the system. When the browser program 38 receives a Web page containing an EMBED tag, a download of the file referenced by the tag is commenced, and the file type is analyzed. If the file is of a type handled by the player program 36, e.g. MPEG, the browser program 38 initiates
10 the player program 36. The contents of the tag are then passed by the browser program 38 to the player program 36.

 The player program 36 (FIGURE 1) provides the Smart Mirroring services facilitated by the MSP 32. The operation of the player program 36 is shown in detail in FIGURE 3.

15 The player program first analyzes the EMBED tag to determine if there is an "SM" (Smart Mirror) parameter (step 60); the presence of such a parameter indicates that the embedded clip is enabled for Smart Mirroring. Data associated with the "SM" parameter specifies the particular content provider from which the desired clip originated, as well as the group of mirror servers that particular
20 content provider uses.

 If the player program 36 determines that the EMBED tag references a video clip or other content handled by the system (step 62), the transfer of the embedded clip from the content provider 22 is stopped. The player program 36 then extracts access control or rating information from the EMBED statement
25 (step 64), if any exists. This rating information is compared against the reference levels stored in the configuration file stored at the user terminal 12 (step 66). If rating information does not exist for the clip, the configuration file is queried to determine whether unrated clips, as defined below, may be played (step 68). Based on the foregoing information, the player program 36 will
30 authorize or decline the viewing of the desired clip.

 If playback is authorized, the player program 36 attempts to find the referenced clip on the local computer belonging to the user terminal 12 (Step

70). If it exists there, it is not re-downloaded, and can be played directly on the computer (from the disk or from RAM) (step 72). However, the time and date of creation of the clip on the local computer is first verified against the time and date for the clip available on the network, to determine if the stored clip is the most recent version (step 74). If it is not, the stored clip is discarded (step 76) and the download proceeds as follows.

If the clip does not exist on the local computer, the player creates a new URL (step 78) in the following form: "http://", plus the IP address of the selected Smart Mirror site stored in the configuration file, plus the path name to mirror files (e.g. "/pub/mirror/"), plus the name of the content provider taken from the "SM" parameter in the EMBED statement, plus the filename taken from the EMBED statement. The constructed URL is used to retrieve the selected clip from the appropriate Smart Mirror site selected by the configuration utility (step 80). If more than one set of delivery sites exists for different content providers, the "SM" parameter is further used by the player program to determine which Smart Mirror site in the configuration file is to be used in the constructed URL (step 82). In one embodiment of the invention, site selection is performed at least partially by a redirection server. This embodiment will be described in detail below with reference to Fig. 4.

If the clip corresponding to the constructed URL is not found at the Smart Mirror site, or is unable to be accessed, then the download proceeds from the next-highest ranked Smart Mirror site in the configuration file (step 84), or the next-highest ranked delivery site selected by the redirection server (see Fig. 4). If all delivery sites fail, the download proceeds from the original content provider's site as specified directly by the EMBED statement.

If playback is disallowed, the player prevents the clip from being transferred (step 88) and displays a bitmap (step 90) advising the user that the download is not be permitted.

If the player program 36 determines that the EMBED tag references a video clip or other content not handled by the system, the player will check whether the access control level set in the configuration file allows the user to play these clips or other files which are considered "unrated" (step 92). If so,

the clip is transferred from its original content provider 22 by traditional means (step 94), and the player program 36 displays the downloaded file (step 96). If not, the player prevents the clip from being transferred (step 98) and displays a message (step 100) advising the user that the download is not permitted.

5 Upon download, the data file representing the desired clip is stored within a specified data area on the local computer, usually on the hard disk, belonging to the user terminal 12 (step 102). In one embodiment, this data area can be managed on a least-recently-used basis by the player program 36. That is, if no room in the data area remains for a new clip, the least-recently-used (or viewed)
10 clip or clips can be discarded to make room (step 104).

In one embodiment of the invention, the client program 36 is capable of sending messages to the MSP 32 (step 106) to reflect whether downloads were successful. This message contains the Internet address of the user terminal 12, the identity of the selected server set, the Internet address of the site used to
15 accomplish the download, the Internet addresses of all sites which failed, the name of the file downloaded, and the time to download the file. This information can also be used by the MSP 32 to track file downloads and to determine, in real time, whether there are any problems with any Smart Mirror sites.

20 Alternatively, the client program 36 can maintain a small local database of file transfer performance. Each download would then be timed. Specifically, information can be gathered on the time it takes a server to begin sending the requested file, the stability of the data transfer rate, and the error rate of the transfer. At some interval (e.g. weekly or once every 100 downloads), a
25 message containing the accumulated file transfer performance information, as well as the user and server information discussed above, would be sent (automatically or upon request) to the MSP 32 (step 106) to update the MSP's database. This additional information increases the MSP's "knowledge" of network performance without incurring any additional testing overhead.

30 This data is especially valuable in ascertaining the performance of delivery sites, for the purpose of assessing the quality of service purchased from the delivery site provider, and for documenting the quality of service to content

providers, to support the cost of the system. It is recognized, however, that much of the same information can be obtained through new users running the configuration utility 34.

In one embodiment of the invention, information derived from the database of the MSP 32 can be used to predict an improved delivery site for a given user without actually having to run the configuration utility to test the network between that user and a list of potential delivery sites. More specifically, the aggregate network performance data contained in the MSP 32 database is analyzed in terms of the performance differences between a given Internet IP address and a number of different delivery sites. Based on this analysis, a correlation can be drawn between a user's IP address and a delivery site that offers better data delivery performance. This correlated data is used to produce a look-up table which can be utilized by the MSP 32 in the process of delivery site selection.

In practice, a redirection server is included within the MSP network. In one embodiment, functions of the redirection server are implemented within and performed by the MSP 32 server. In alternative embodiments, several redirection servers may be employed at various locations throughout the network.

In operation, the redirection server will acquire the IP address of the user when the user requests a file that is managed by the MSP delivery system. A request is made of the redirection server through an EMBED statement, as discussed above. The EMBED statement can explicitly refer to a specific server IP address, identifying a single redirection server (e.g. the MSP 32), or can contain a script or executable program code (e.g. in JavaScript or the Java programming language) specifying a plurality of redirection servers, to which access can be attempted in sequence or in a random order. The user's IP address is then determined by the server through traditional means, as HTTP (HyperText Transport Protocol) requests typically include information on the requester's address. The redirection server then maps the user's IP address to an optimum delivery site located in the look-up table and forwards the delivery site address to the user. The user's client program 36 or the redirection server

then redirects the file request so that the file is delivered from the optimum delivery site.

An important factor in designing the look-up table derives from the large number of addresses on the Internet. Every computer connected to the Internet is assigned an address. In the conventional Internet addressing scheme, An address is comprised of a 4-byte value that, by convention, is expressed by converting each byte into a decimal number (0-255) and separating the bytes with a period. For example, an address for the "www.intervu.net" server is 192.215.147.185. With this addressing scheme, there are over four billion possible addresses on the Internet. Given the current state of the art, a look-up table including four billion addresses is too large to be useful in the embodiment described herein, so a way was sought to compress the table.

It has been found that the selection of an optimum delivery site can often be correlated to the first byte of a user's IP address. This correlation is such that, after running the configuration utility 34, a statistically significant number of users having the same first-byte address would select the same delivery site, or another site from a small group of delivery sites. It has been observed that these delivery sites, in comparison to uncorrelated delivery sites, provide improved performance for most users having the same first-byte address. The correlation is significant enough that in one embodiment of the invention, a compressed look-up table is formed comprising a list of first-byte IP addresses numbering 0-255, and for each address, a list of delivery sites providing improved performance for users having corresponding IP addresses. In this way, one or more delivery sites are mapped to each entry in the address list.

It should be noted that in certain circumstances, some of the first-byte IP addresses in the look-up table might not be mapped to a corresponding server. This would be the case when, for example, few users have IP addresses corresponding to a particular entry in the table, and not enough of them have run the configuration utility 34 to generate statistically significant or reliable results. In this case, one or more default servers can be specified in the database for use when a specific server is not identified in the look-up table.

In a preferred embodiment, the look-up list is stored at a single redirection

server, e.g. the MSP 32. However, the look-up list and related programming could also be stored at a web page server, a content provider 22 server or a combination of the above.

The redirection server is accessed to select a delivery site (step 82, Fig. 3) as follows, and as described in Fig. 4. The user, via the user terminal browser, requests a file referenced on a web page by clicking on a link having an EMBED statement for that file. If the file request is for a file managed by the MSP's delivery system, the file request is forwarded to the redirection server (e.g. the MSP 32) using conventional HTML file request semantics, i.e. server "GET" (step 120).

The redirection server examines the incoming request and determines the user's network address (IP address) using the "REMOTE_HOST" variable supplied by the web server (step 122). It should be noted that the embodiment of the invention described in Fig. 4 can be accomplished without special-purpose client program (such as the client program 36) installed on the user terminal 12. If no client program 36 is installed (step 124), and the user, when queried by a script or downloaded program, does not want to install a client program (step 126), then site selection will be performed entirely by the redirection server.

The redirection server then analyzes the user's IP address and examines the list of potential delivery sites on the look-up table to determine what delivery site or sites are correlated with the user's IP address (step 128). The redirection server then chooses a delivery site address, if more than one address corresponds to the user's look-up table entry (step 130). If a single delivery site address is found, then an HTTP redirection response is used to deliver the file from the selected delivery site (step 132). According to the HTTP specification, a file requested from a certain server can be redirected by that server to another location without any user or client program intervention. Site selection is then complete (step 134). To complete the transaction, the client program requests a copy of the file from the selected delivery site and the delivery site server retrieves the file and delivers it to the user.

If the user wants to install the client program 36 (step 126), then the program will be downloaded and installed (step 136) by traditional means. If the

client program 36 was already installed (step 124), or if installation has just completed, the file request proceeds differently, and is able to utilize both server-side and client-side processing to retrieve a file in the most efficient manner possible.

5 First, the redirection server analyzes the user's IP address and examines the list of potential delivery sites on the look-up table to determine what delivery site or sites are correlated with the user's IP address (step 138). In some cases, it is possible that the redirection server storing the look-up table will generate (step 140) and send (step 142) to the user terminal 12 a small file including a
10 sublist of delivery sites and rely on the client program 36 to make the final delivery site selection (step 144). There are a number of reasons this may occur.

First, as previously noted, an analysis of aggregate network performance data may indicate that a number of delivery sites (as opposed to a single delivery
15 site) offer improved performance to certain users having the same first-byte address range. This occurs, for example, where a number of users, all having the same first-byte address or range of addresses within a single first-byte address, selected different delivery sites after running the configuration utility. In this case, the client program 36 can take the sublist downloaded by the
20 redirection server and compare it to a saved list of mirror sites which had been previously selected by running the configuration utility (see Fig. 3, step 54). If a match is found, the client program redirects the file request to the matching delivery site (step 146). The client would also have the option to ignore the selection make by the client program and make a best-guess selection from
25 among the delivery sites contained in the look-up list, since one or more of these delivery sites may offer improved performance over any delivery site previously selected by the configuration utility, especially if the configuration utility has not been run recently.

Second, a situation may also arise where the number of delivery sites and
30 IP addresses managed by the MSP has increased to a point where the look-up table becomes too large to be practical, i.e., the redirection server is delayed in responding to a new file request because it is busy searching through the look-

up tables on behalf of a previous file request. If this is the case, the task of searching through the look-up table is split between the redirection server and the client program 36. Thus the user terminal 12 is required to do some of the processing and also act as a router by making the final delivery site selection.

5 In order to accomplish this, the redirection server subdivides the look-up table into smaller sublists with a given range of addresses. Thus, when the redirection server receives a request for a file, the server will map the user's IP address (step 138) to the sublist with the corresponding address range, and then generate (step 140) and download (step 142) to the user a small file containing
10 the sublist. The client program 36 then acts as a router, selecting a delivery site (step 144) from among the delivery sites in the sublist look-up table, by either mapping the user's IP address to the proper correlated delivery site address on the sublist, or by looking for a match between a delivery site on the sublist and a delivery site on the prioritized list of sites saved by the configuration utility 34
15 after previously running network performance tests, as described above. Once the client program has selected a delivery site, the client program redirects the file request to the selected delivery site (step 146). By splitting the work between the redirection server and the user terminal 12, the file request can be processed quicker, thereby reducing the delay between the time the user
20 requests the file and the time when the file is received by the user terminal 12.

Although the MSP 32 redirection look-up table will be updated frequently in response to changes in network performance, the sublists passed down to the user are expected to be satisfactory for at least a few days. Thus, the client program 36 can reuse the list for a few days before acquiring a new sublist.
25 This function can be implemented, for example, by a script embedded in the web page hosting the file request statement. The script can query the client program 36 for an expiration date encoded in the sublist stored at the user terminal 12. The expiration date and sublist type can be passed back to the redirection server and if necessary, the server will pass down a new sublist.

30 It is possible that a particular user may request a file managed by the MSP delivery system, and the system will have no knowledge of that user (i.e., no matches are found between the user's IP address and delivery sites in the look-

up table). In this case, the redirection server will select a delivery site for the user (steps 130 and 144) from a default list stored at the redirection server. When the user downloads the requested file from this delivery site, he may be prompted via an embedded script or program to acquire the configuration utility file from the MSP 32 for the purpose of improving content delivery, or will be temporarily assigned a server which is an approximation (e.g. has a similar, but not identical, first-byte IP address) for improved content delivery. If the user acquires and runs the configuration utility, the user terminal 12 will begin providing the MSP 32 with the results of various network tests, as described above, and this information will ultimately become part of the MSP database of network performance data. Thus, on the next file request from that new user, the redirection server will be able to map the user's IP address to a more appropriate server based on an analysis of the additional network performance data.

By modifying the look-up table, the redirection server can also perform load balancing and management with respect to file requests. If the MSP 32 (or an individual controlling the MSP 32) determines in advance that sections of the network will be down for a period of time, or if certain delivery sites must be shut down for a period of time, the look-up table can be modified so that the redirection server will point users to alternative delivery sites. Thus, it is possible that on two back-to-back requests for the same file, a user may be directed to two different delivery sites.

It should be appreciated that the file comprising a sublist of the look-up table can function as the delivery site file which is used by the configuration utility 34 to run network tests. The sublist represents delivery sites which have already been screened for improved performance via aggregate network performance data derived from a group of other users who have previously conducted network tests. This sublist of delivery sites can be further pared-down or prioritized by running the configuration utility 34 against the list and performing the sequence of network tests retained from a previous delivery site file. Once the sublist is prioritized, the configuration utility 34 saves the list to the (encrypted) configuration file as described above.

In the following exemplary description, the Smart Mirror system utilizes the look-up table to locate an improved delivery site from which to download a video clip referenced on a Web page advertising banner. When dealing with relatively unpopular content such as advertising, it is important to deliver advertising content to the intended customer as quickly as possible, so the potential customer does not lose interest. With respect to the content in a Web page advertising banner, the same is true: it is desirable to be able to deliver the video to the customer as quickly as possible without having to download a substantial amount of software in advance. Therefore, as an alternative to utilizing the invention as a premium service, a service provider can distribute advertising banners on Web pages across the Internet. Each banner would reference a video clip, and all of the clips would be stored at a content provider's server established for that purpose.

In this example, the client program 36 is referred to as a "player program," and it has the added functionality of retrieving and playing back video data as previously described. A browser 38 (such as Netscape Navigator or Microsoft Internet Explorer) is used to view Web content at the user terminal 12 and to communicate between the user terminal 12 and other computers on the Internet 10.

When the client selects an advertising banner displayed on a Web page, an EMBED tag encoded in the banner is examined by the browser. The browser commences a download of the file referenced by the EMBED tag and the file type is analyzed. If the file is of a type handled by the player program, e.g. MPEG, the browser initiates the player program by passing the tag to the player program. The player program examines the EMBED tag for the file name of the video clip and attempts to find the clip on the local computer belonging to the user terminal. If it exists there, it is not re-downloaded, and can be played directly on the computer (from the disk or from RAM). If the video clip is not on the local computer, the player program analyzes the EMBED tag to determine if there is an "SM" (Smart Mirror) parameter; the presence of such a parameter indicates that the embedded clip is enabled for Smart Mirroring. The SM parameter further specifies the address of the redirection server, and as

discussed above, the particular content provider from which the desired clip originated, as well as the group of mirror servers that particular content provider uses.

If the player program recognizes a parameter specifying a redirection server, the player program will invoke the browser to request the referenced ad banner file from the redirection server via conventional HTML file request techniques. The redirection server examines the incoming request and determine the IP address of the client. The redirection server maps the client's IP address to the look-up table and retrieve a delivery site address or a sublist. The delivery site address (or a sublist of delivery sites and IP addresses) is returned to the player program by the redirection server. If a sublist is returned, the player program selects a single delivery site from among those listed, as previously described. Once a single delivery site is determined, the player program creates a new URL in the following form: "http://", plus the IP address of the selected delivery site, plus the path name to mirror files (e.g. "/pub/mirror/"), plus the name of the content provider taken from the "SM" parameter in the EMBED statement, plus the filename taken from the EMBED statement. The constructed URL is used to retrieve the selected clip from the delivery site selected from the look-up table.

If the SM parameters in the EMBED tag do not specify a redirection server, or if the redirection server is unavailable and does not return a response, the player program will select an IP address from the prioritized list of Smart Mirror sites stored in the configuration utility file and substitute this address in the URL. If the player program does not detect the presence of the prioritized list of smart mirror sites (for example, if the user never downloaded the configuration utility 34), then the player program will request the video using the address of the original content provider's server, as specified directly by the EMBED statement.

In applications not specifically targeted to the delivery of advertising content, the provision of download information to the MSP facilitates the use of the invention as a premium subscription-based service. As successful downloads are tracked in a database, each user can have an associated "account" to track charges. The user can be charged for use of the Smart

Mirror system by the file, by the megabyte, by the month, or by other known means. In one embodiment, the EMBED tag associated with a file contains billing information, or a "price" for the file. The invention's tracking of download performance allows discounts or credits to be issued if downloads are found to
5 be unduly difficult or slow.

To ensure that files stored on Smart Mirror delivery sites are used only by authorized users of the invention (e.g. those users having paid accounts), the files stored at the delivery sites can optionally be in encrypted form, and the downloading step described above can include a decryption step. Such
10 encryption and decryption can be performed by well known means.

As discussed above, the clips managed by the invention can have content rating information associated therewith. This is accomplished by providing a "PG" parameter in the EMBED statement corresponding to the clip. In one embodiment, four characteristics are rated: nudity, sexuality, profanity, and
15 violence. Accordingly, the PG parameter can be specified by a four-digit argument. Each characteristic is rated on a scale of one to three. One corresponds to no filtering (i.e. all content is allowable), two corresponds to some filtering (e.g. equal to levels typically allowed in broadcast television), and three corresponds to the most extensive filtering (e.g. for children). The ratings
20 levels contained in the EMBED statement for a file are compared to the ratings filter levels contained in the configuration file stored at the user terminal 12 in the foregoing authorization process, and only authorized files are transferred.

In view of the above, it will be appreciated that embodiments of the invention may be employed in many different applications to permit the
25 acquisition and analysis of performance data for networks between a given user and content provider or delivery site. Thus, although the described embodiment illustrates the system operating within the context of the Internet, and with an Internet-type addressing scheme, it is recognized that such a system could prove to be useful in other network environments, including but not limited to
30 corporate "intranets."

Moreover, although the illustrative embodiments are described primarily for use in a video delivery system, it should be recognized that a system

according to the invention can be used to distribute various other kinds of computer data (e.g. application programs, database files and other business information, virtual reality files, multimedia such as Macromedia Shockwave files, and large text files such as books) as well. Such other types of data can
5 be managed by the invention in different content provider groups as discussed in detail above; a different type of program (rather than the player program 36) typically would be invoked at the user terminal 12 to view or use other types of data.

It should also be noted that certain functionality described as performed
10 at the user terminal 12 (specifically, certain functions performed by the configuration utility 34, or client/player program 36) can be implemented as a standalone program, as a "plug-in" or "helper application" to run within a browser program, or as a Java applet downloaded from a delivery site to run within a browser environment. For user terminals capable of running the
15 Microsoft Windows operating system, an environment known as Microsoft ActiveX is also useful.

While certain exemplary structures and operations have been described, the invention is not so limited, and its scope is to be determined according to the claims set forth below.

What is claimed is:

1 1. A system for selecting a delivery site, located on a distributed
2 computer network, from which to receive a file, the system comprising:
3 a communications network;
4 at least one content provider connected to the network, wherein
5 the content provider stores at least one file;
6 at least two delivery sites connected to the network, wherein each
7 delivery site stores the same file as the content provider;
8 a redirection server connected to the network;
9 a user terminal connected to the network;
10 a network testing apparatus for testing the network and selecting
11 a preferred delivery site from all delivery sites connected to the network; and
12 a downloading apparatus for downloading the file from the preferred
13 delivery site specified by the redirection server to the user terminal.

1 2. A system according to claim 1, wherein the network testing apparatus
2 comprises a first storage medium connected to the user terminal, wherein the
3 first storage medium embodies a first software program for testing the network
4 and selecting a preferred delivery site from all delivery sites connected to the
5 network.

1 3. A system according to claim 2, wherein the downloading apparatus
2 comprises a second storage medium connected to the user terminal, wherein the
3 second storage medium embodies a second software program capable of
4 downloading the file from the preferred delivery site to the user terminal.

1 4. A system according to claim 3, wherein the redirection server
2 comprises a database.

1 5. A system according to claim 4, wherein the first software program
2 uses a list of delivery sites created by the database.

1 6. A system according to claim 5, wherein the first software program

2 uses results from a network analysis procedure to choose the preferred delivery
3 site from the list of delivery sites.

1 7. A system according to claim 5, wherein the first software program
2 uses results from a network analysis procedure to create a prioritized list of
3 preferred delivery sites.

1 8. A system according to claim 6, wherein the network analysis
2 procedure is performed at least in part at the user terminal.

1 9. A system according to claim 8, wherein the network analysis
2 procedure comprises a plurality of network tests.

1 10. A system according to claim 9, wherein the network tests are
2 selected from a list of tests maintained in the database.

1 11. A system according to claim 9, wherein results from the network
2 tests are provided to the database.

1 12. A system according to claim 9, wherein the first software program
2 performs the network tests in conjunction with the user terminal.

1 13. A system according to claim 11, wherein the results are stored in the
2 database.

1 14. A system according to claim 13, wherein the database maps an
2 address for the user terminal to at least one delivery site.

1 15. A system according to claim 14, wherein the second software
2 program performs the download in conjunction with the user terminal.

1 16. A system according to claim 15, wherein the second software

2 program performs the download from a server specified by the database.

1 17. A system according to claim 12, wherein the first software program
2 can be executed at any time.

1 18. A system according to claim 17, wherein the first software program
2 is executed automatically before any file is downloaded.

1 19. A system according to claim 1, wherein the file comprises a video
2 clip.

1 20. A system according to claim 7, wherein the second software program
2 uses the prioritized list.

1 21. A system according to claim 20, wherein the second software
2 program accesses the database.

1 22. A system according to claim 21, wherein the second software
2 program first attempts to use a first delivery site specified in the prioritized list.

1 23. A system according to claim 22, wherein if the first delivery site fails,
2 the second software program will attempt to use succeeding delivery sites
3 specified in the prioritized list.

1 24. A system according to claim 23, wherein if all delivery sites in the
2 prioritized list fail, the second software program will attempt to download
3 directly from the content provider.

1 25. A system according to claim 1, wherein the file has characteristic
2 information associated therewith.

1 26. A system according to claim 25, wherein the characteristic
2 information is stored within an "EMBED" statement corresponding to the file.

1 27. A method for determining a set of preferred delivery sites from a
2 plurality of delivery sites on a network, comprising the steps of:
3 obtaining a list of delivery sites;
4 performing at least one network test;
5 processing results from the network test; and
6 mapping the delivery sites to IP addresses according to the test
7 results.

1 28. The method of claim 27, wherein the performing step comprises
2 testing some of the delivery sites in the list.

1 29. The method of claim 28, wherein the performing step comprises:
2 selecting a subset of delivery sites from the list; and
3 running at least one test for each of the delivery sites in the subset.

1 30. The method of claim 29, wherein the test is selected from a plurality
2 of tests.

1 31. The method of claim 30, wherein the plurality of tests is specified in
2 a database connected to the network.

1 32. The method of claim 31, wherein the test is selected from the
2 plurality by way of random selection based on a frequency specified in the
3 database.

1 33. The method of claim 27, further comprising the step of sending the
2 results to a database connected to the network.

1 34. The method of claim 27, wherein the processing step comprises
2 weighting the results.

1 35. The method of claim 27, wherein the mapping step comprises storing
2 the set of delivery sites in a preferred order for each of a plurality of address
3 ranges.

1 36. The method of claim 35, wherein each address range corresponds to
2 all IP addresses having the same first byte.

1 37. The method of claim 36, wherein the preferred order is determined
2 by the test results.

1 38. The method of claim 27, wherein the list of delivery sites is obtained
2 from a database.

1 39. A method for downloading a file from a delivery site to a user
2 terminal, comprising the steps of:
3 identifying a desired file;
4 determining whether the desired file is mirrored on one or more delivery
5 sites;
6 if the desired file is mirrored, finding a preferred delivery site in a list of
7 delivery sites mapped to IP addresses; and
8 downloading the desired file from the preferred delivery site.

1 40. The method of claim 39, wherein the determining step comprises the
2 substeps of querying a redirection server and receiving a response.

1 41. The method of claim 40, wherein the querying step is performed by
2 a script.

1 42. The method of claim 41, wherein the script is embedded within a
2 Web page.

1 43. The method of claim 40, wherein the querying step is performed by
2 a program downloaded to the user terminal.

1 44. The method of claim 40, wherein the retrieving and downloading
2 steps are repeated until the downloading step is successful.

1 45. The method of claim 44, further comprising the step of downloading
2 the desired file from a content provider site if the repeated retrieving and
3 downloading steps are unsuccessful.

1 46. The method of claim 40, wherein the finding step comprises the
2 substeps of:
3 determining an address for the user terminal;
4 comparing the address to a list of addresses in a database;
5 if the database contains sufficient data, locating the preferred
6 delivery site corresponding to the address; and
7 if the database does not contain sufficient data, locating a default
8 delivery site.

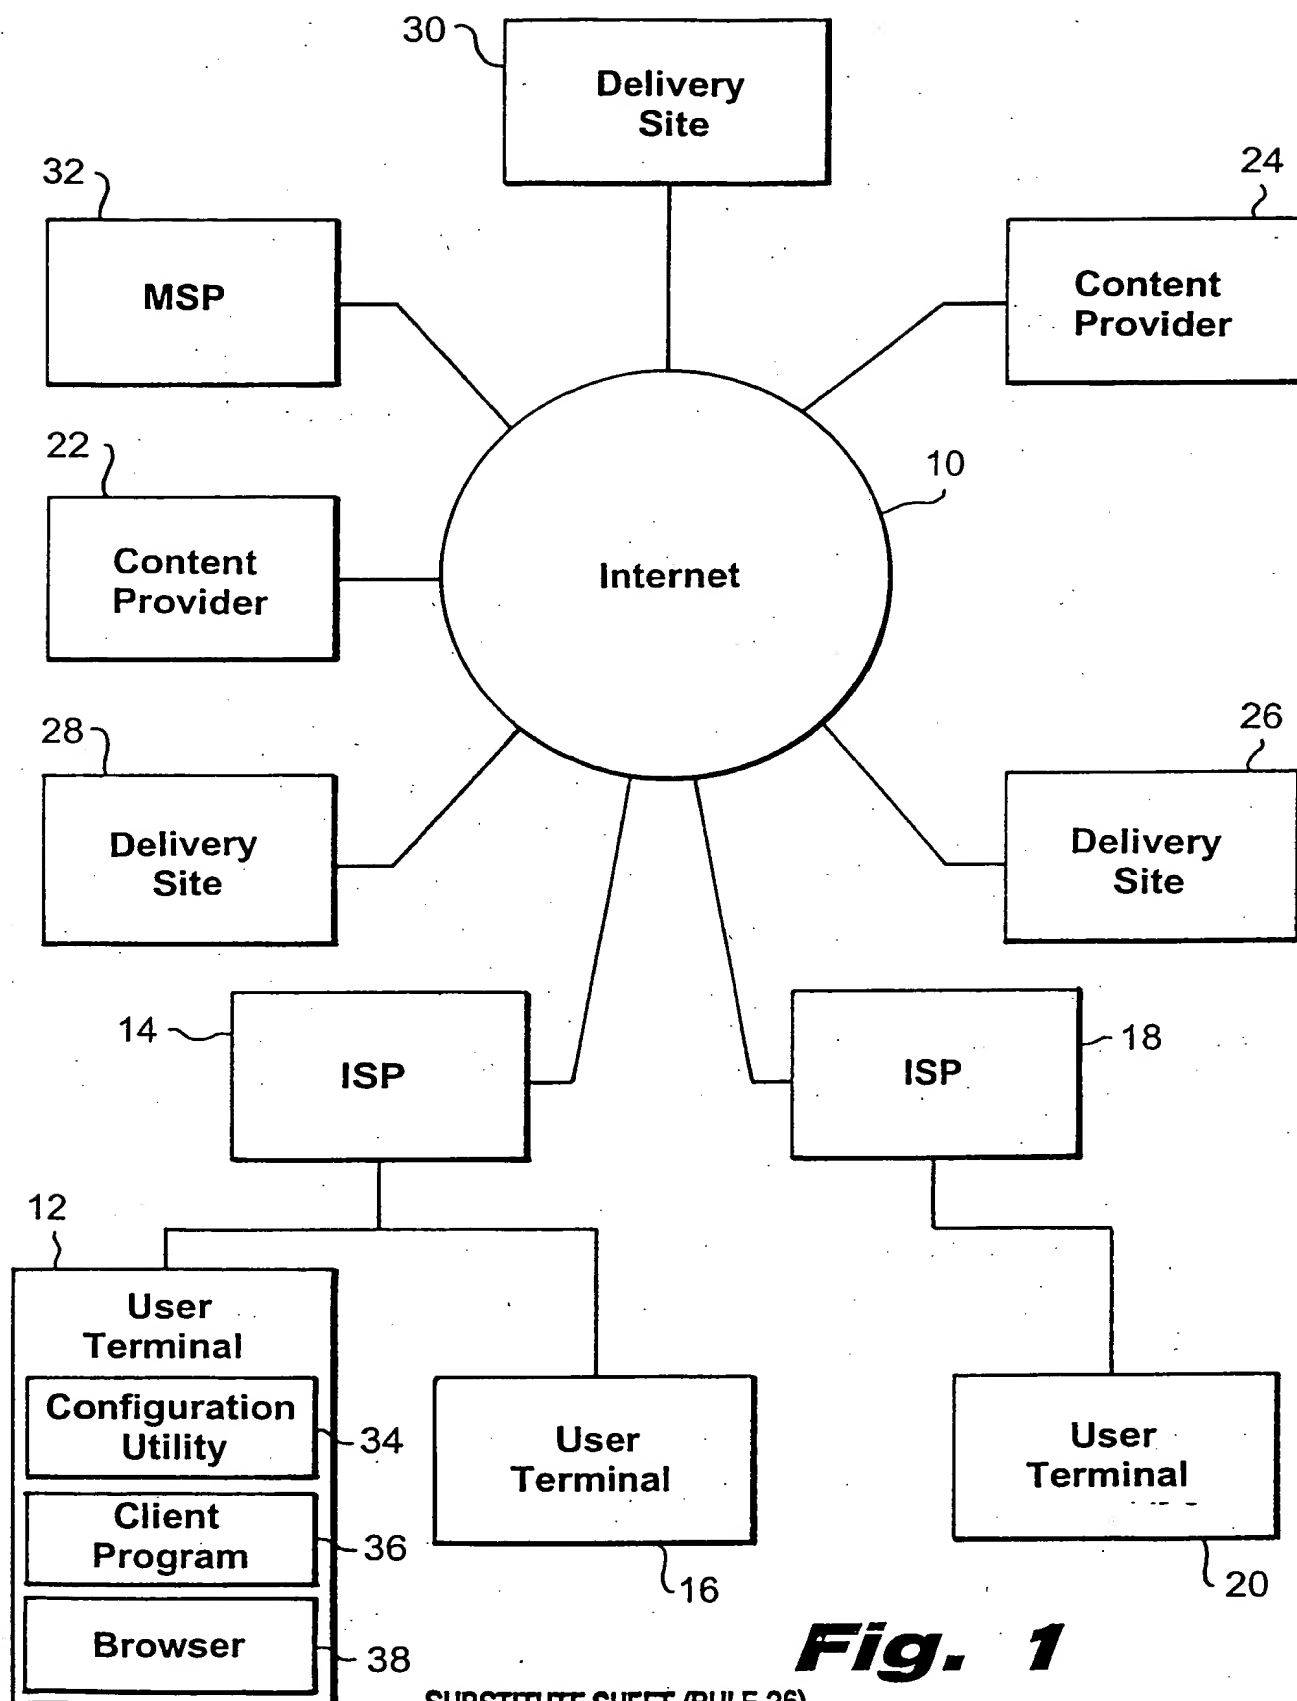
1 47. The method of claim 46, wherein the comparing step is performed
2 at a computer hosting the database.

1 48. The method of claim 47, wherein the locating steps are performed
2 at a computer hosting the database.

1 49. The method of claim 46, wherein the first locating step comprises the
2 substeps of:
3 isolating a sublist of addresses; and
4 selecting the preferred delivery site within the sublist.

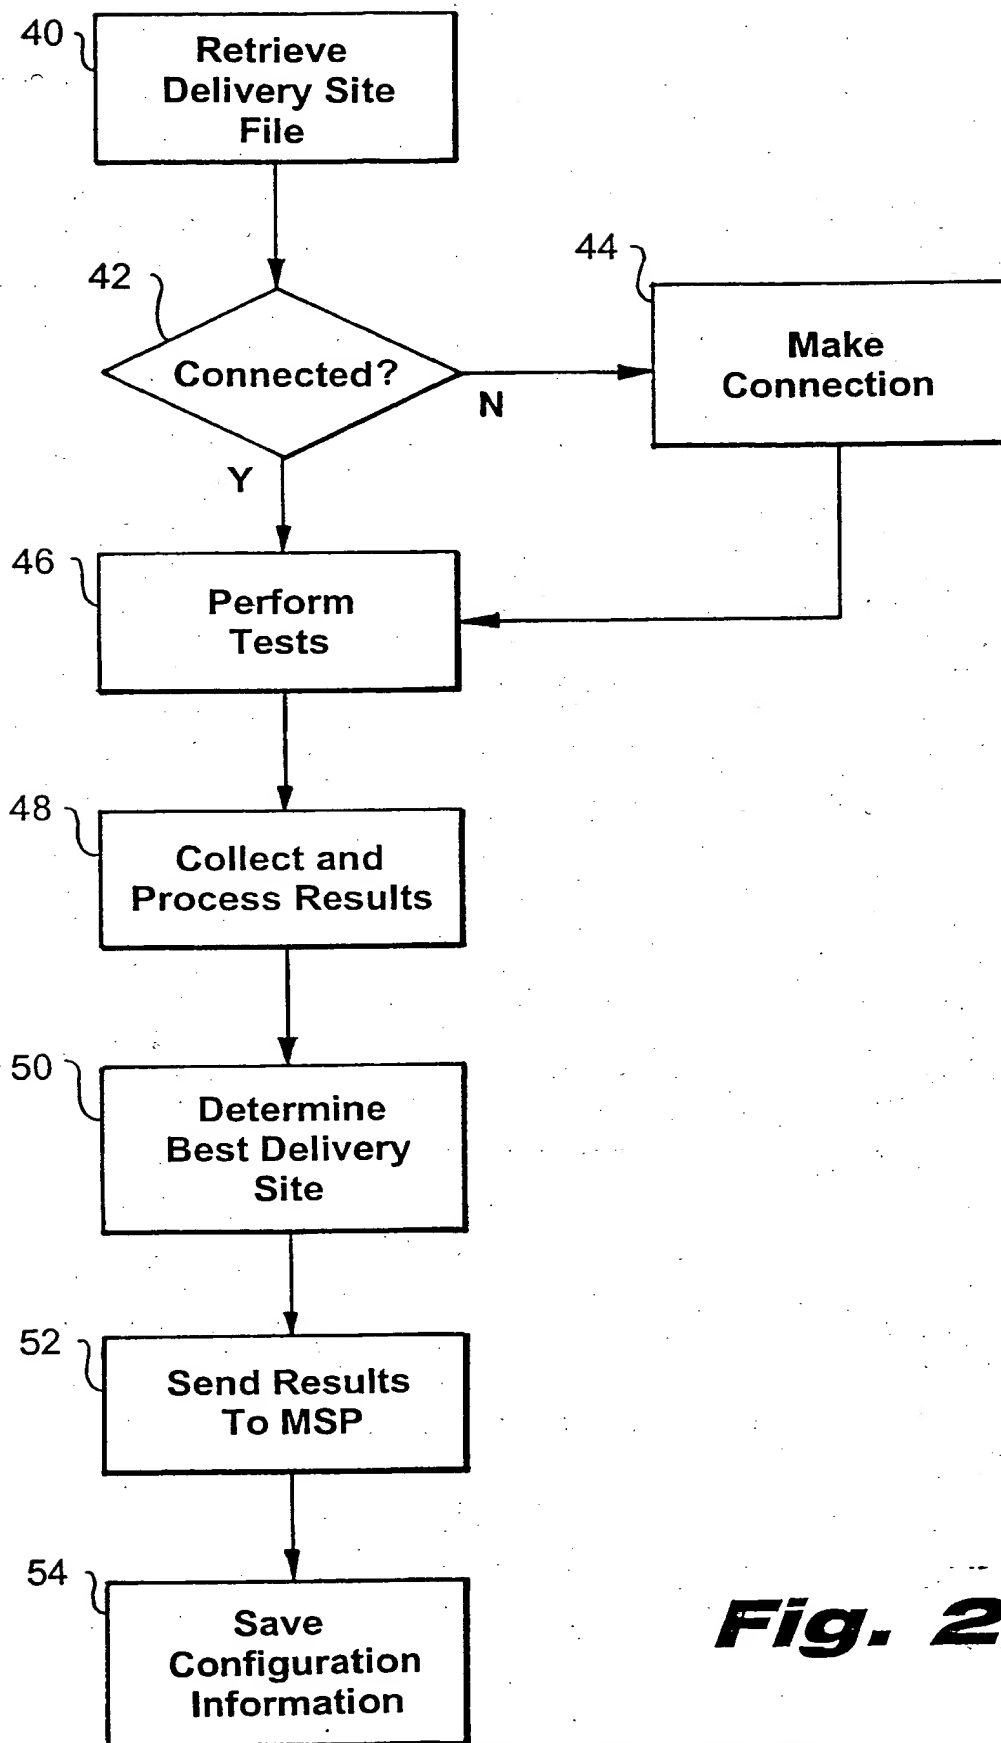
1 50. The method of claim 49, wherein the selecting step is performed at
2 the user terminal.

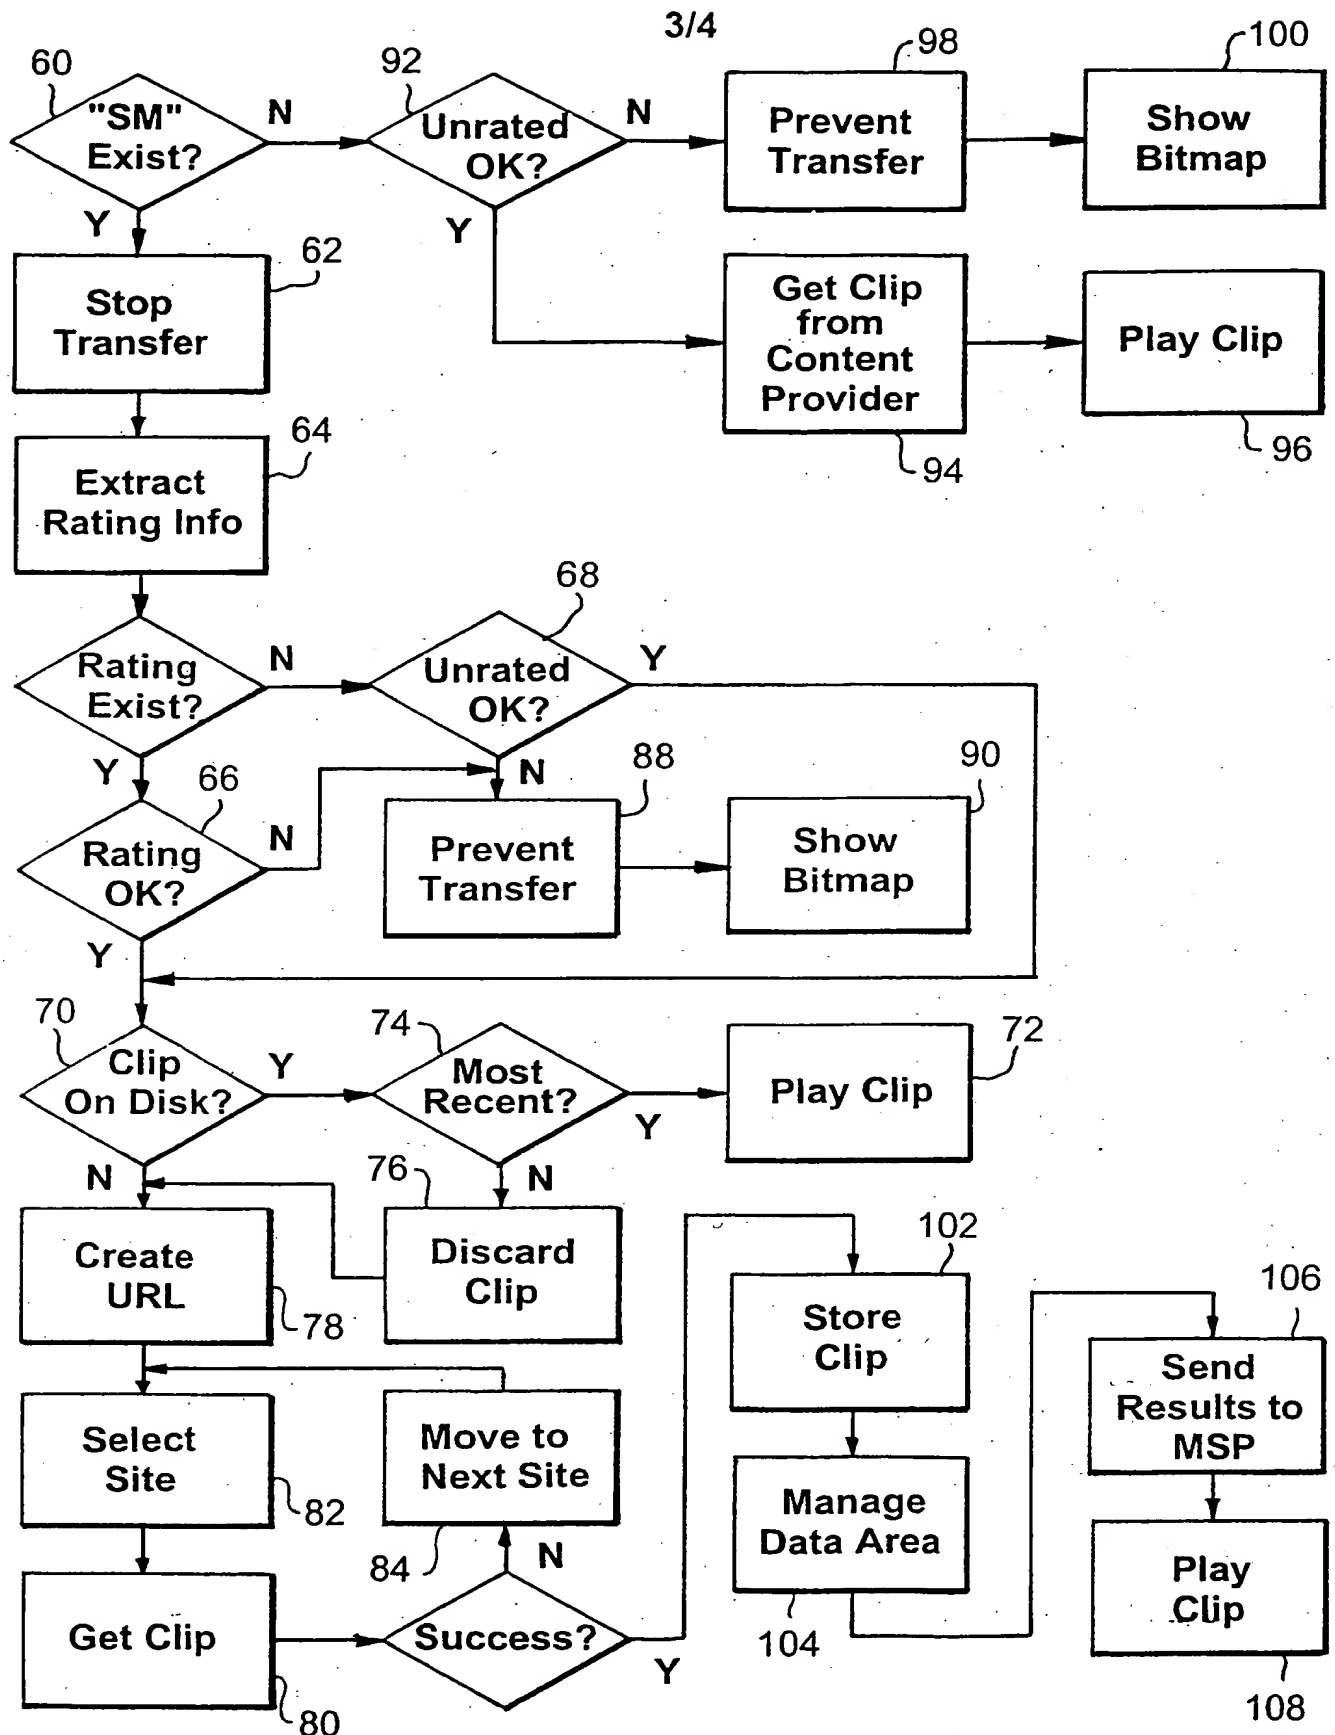
1/4

**Fig. 1**

SUBSTITUTE SHEET (RULE 26)

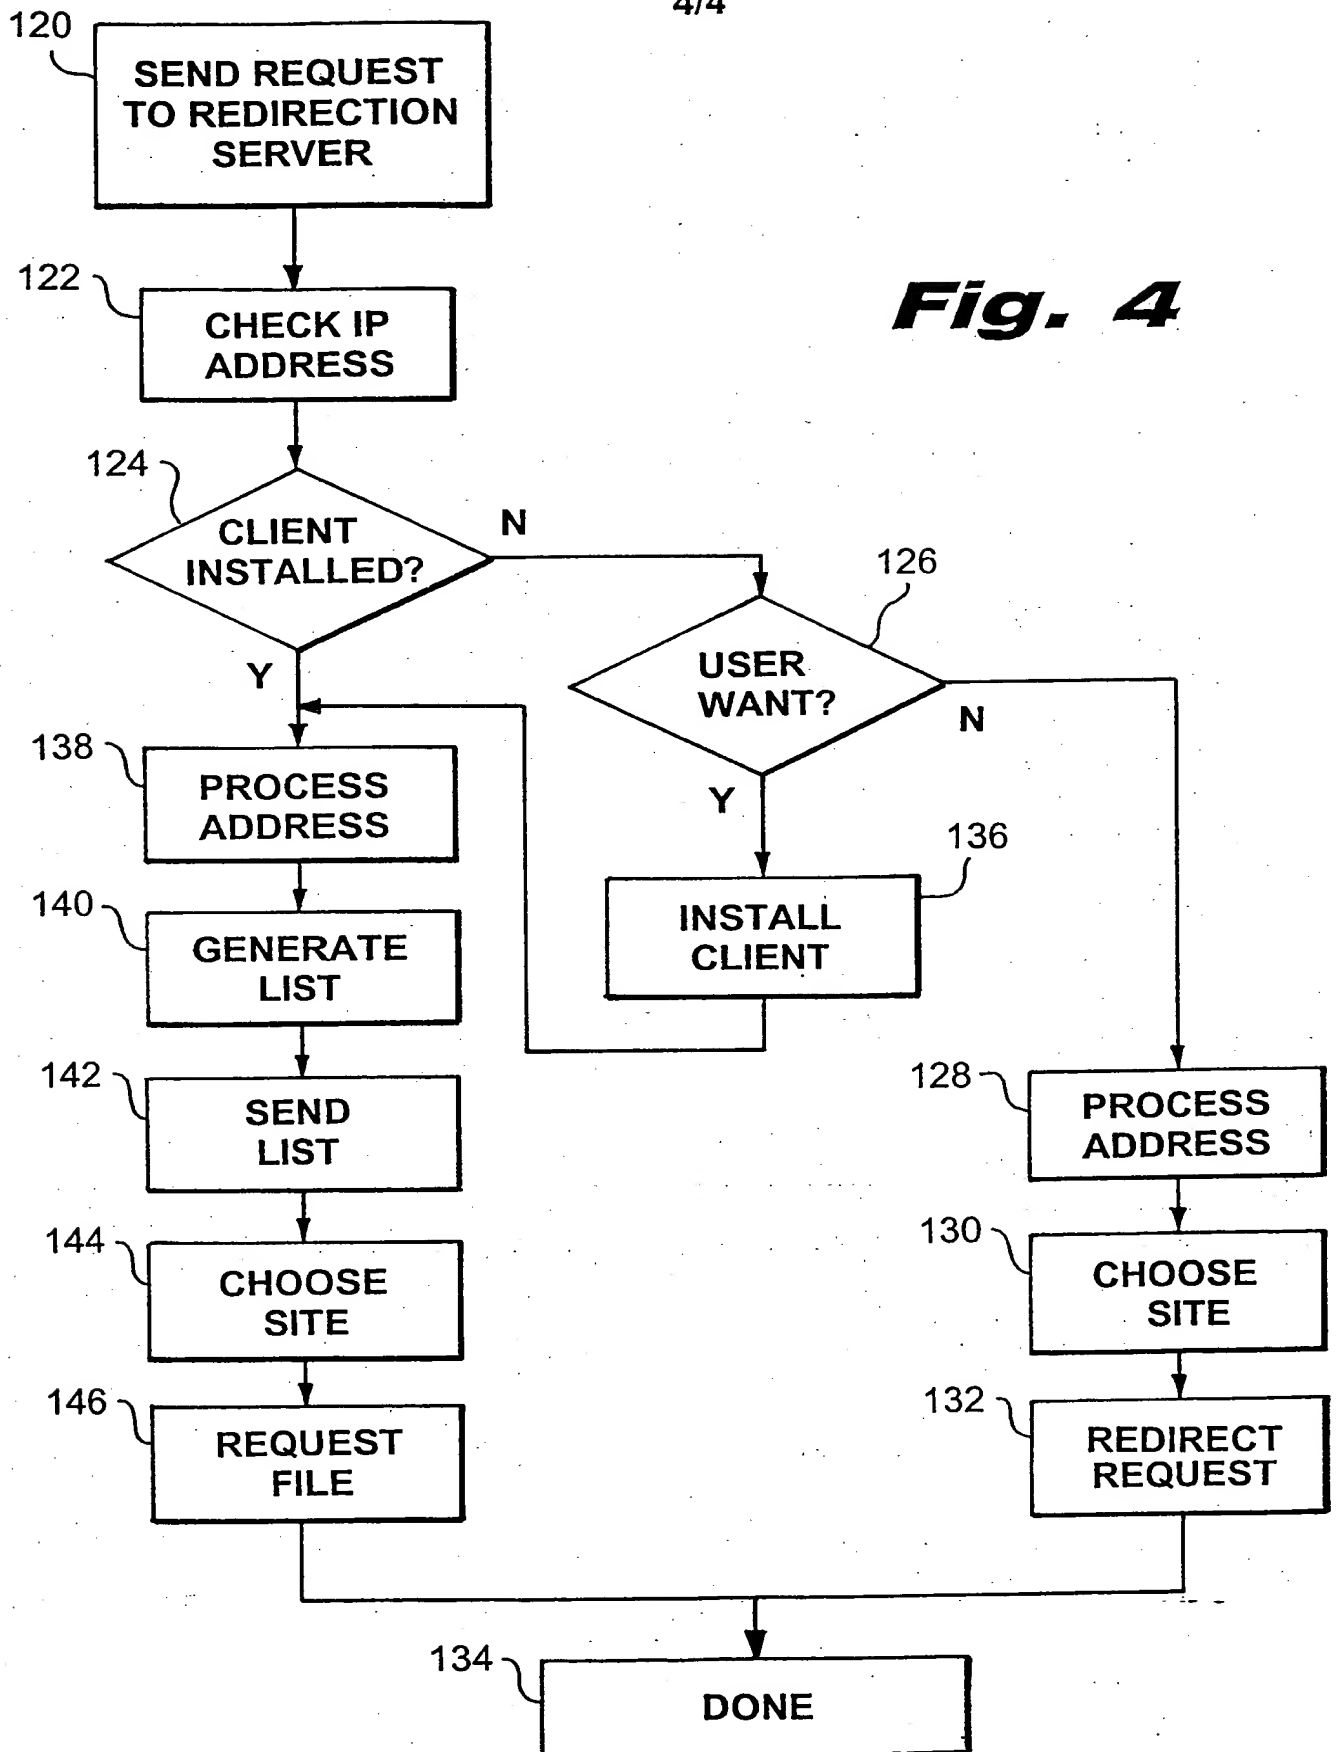
2/4

**Fig. 2**

**Fig. 3**

SUBSTITUTE SHEET (RULE 26)

4/4

Fig. 4

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 98/12784

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 G06F9/46 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	"SINGLE SYSTEM IMAGE AND LOAD BALANCING FOR NETWORK ACCESS TO A LOOSELY COUPLED COMPLEX" IBM TECHNICAL DISCLOSURE BULLETIN, vol. 34, no. 9, 1 February 1992, pages 464-467, XP000301951 see the whole document	1, 27, 39
A	MULLEN T ET AL: "A SIMPLE COMPUTATIONAL MARKET FOR NETWORK INFORMATION SERVICES" PROCEEDINGS OF ICMAS: INTERNATIONAL CONFERENCE ON MULTIAGENT ESYSTEMS, 12 June 1995, pages 283-289, XP002054269 see page 284, right-hand column, line 27 - page 298, left-hand column, line 22 -/--	27

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

21 September 1998

Date of mailing of the international search report

28/09/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Deane, E

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 98/12784

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 94 25919 A (EMC CORP) 10 November 1994 see page 6, line 26 - page 10, line 14; figure 1. ---	27
A	ABBASOV A M: "OPTIMIZING THE LOCATION OF DATABASES WITH COPIES IN COMPUTER NETWORKS" AUTOMATIC CONTROL AND COMPUTER SCIENCES, vol. 22, no. 4, 1 January 1988, pages 58-62, XP000120119 see the whole document ---	1,27,39
A	EP 0 384 339 A (DIGITAL EQUIPMENT CORP) 29 August 1990 see column 5, line 10 - column 7, line 25; figures 1,2 -----	1,27,39

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 98/12784

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9425919 A	10-11-1994	US 5544347 A	06-08-1996
		AU 684773 B	08-01-1998
		AU 6638094 A	21-11-1994
		EP 0695443 A	07-02-1996
		JP 2790378 B	27-08-1998
		JP 8509565 T	08-10-1996
		US 5742792 A	21-04-1998
		US 5664144 A	02-09-1997
EP 0384339 A	29-08-1990	AT 151183 T	15-04-1997
		AU 611605 B	13-06-1991
		AU 4996190 A	13-09-1990
		AU 630291 B	22-10-1992
		AU 7603391 A	15-08-1991
		CA 2010762 A	24-08-1990
		DE 69030340 D	07-05-1997
		DE 69030340 T	20-11-1997
		JP 3116262 A	17-05-1991
		US 5341477 A	23-08-1994

THIS PAGE BLANK (USPTO)